arXiv:1404.6724v2 [cs.DS] 1 May 2014

Approximately Minwise Independence with Twisted Tabulation

Søren Dahlgaard* Mikkel Thorup*

University of Copenhagen,
{soerend,mthorup}@di.ku.dk

May 2, 2014

Abstract

A random hash function h is ε-minwise if for any set S, $|S| = n$, and element $x \in S$, $\Pr[h(x) = \min h(S)] = (1 \pm \varepsilon)/n$. Minwise hash functions with low bias ε have widespread applications within similarity estimation.

Hashing from a universe $[u]$, the twisted tabulation hashing of Pătraşcu and Thorup [SODA'13] makes $c = O(1)$ lookups in tables of size $u^{1/c}$. Twisted tabulation was invented to get good concentration for hashing based sampling. Here we show that twisted tabulation yields $\tilde{O}(1/u^{1/c})$-minwise hashing.

In the classic independence paradigm of Wegman and Carter [FOCS'79] $\tilde{O}(1/u^{1/c})$-minwise hashing requires $\Omega(\log u)$-independence [Indyk SODA'99]. Pătraşcu and Thorup [STOC'11] had shown that simple tabulation, using same space and lookups yields $\tilde{O}(1/n^{1/c})$-minwise independence, which is good for large sets, but useless for small sets. Our analysis uses some of the same methods, but is much cleaner bypassing a complicated induction argument.

1 Introduction

The concept of minwise hashing (or the "MinHash algorithm" according to [1]) is a basic algorithmic tool suggested by Broder et al. [3, 5] for problems related to set similarity and containment. After the initial application of this algorithm in the early AltaVista search engine to detecting and clustering similar documents, the scheme has reappeared in numerous other applications[1] and is now a standard tool in data mining where it is used for estimating similarity [5, 3, 6], rarity [8], document duplicate detection [4, 14, 24, 10], large-scale learning [13], etc. [1, 2, 7, 18].

*Research partly supported by Thorup's Advanced Grant from the Danish Council for Independent Research under the Sapere Aude research carrier programme.

[1] See http://en.wikipedia.org/wiki/MinHash

The basic motivation of minwise independence is to use hashing to select an element from a set S. With a hash function h, we simply pick the element $x \in S$ with the minimum hash value. If the hash function is fully random and no two keys get the same hash, then x is uniformly distributed in S.

A nice aspect of minwise selection is that $\min h(A \cup B) = \min\{\min h(A), \min h(B)\}$. This makes it easy, e.g., to select a random leader in many distributed settings. It also implies that that $\min h(A \cup B) \in h(A \cap B) \iff \min h(A) = \min h(B)$. Therefore, if h is fully random and collision free,

$$\Pr_h[\min h(A) = \min h(B)] = \frac{|A \cap B|}{|A \cup B|} \ .$$

Thus, if we, for two sets A and B, have stored $\min h(A)$ and $\min h(B)$, then we can use $[\min h(A) = \min h(B)]$[2] as an unbiased estimator for the Jaccard similarity $|A \cap B|/|A \cup B|$.

Unfortunately, we cannot realistically implement perfect minwise hash functions where each $x \in S$ has probability $1/|S|$ of being the unique minimum [5]. More precisely, to handle any subset S of a universe $\mathcal{U}$, we need a random permutation $h : \mathcal{U} \to \mathcal{U}$ represented using $\Theta(|\mathcal{U}|)$ bits.

Instead we settle for a bias ε. Formally, a random hash function $h : \mathcal{U} \to \mathcal{R}$ from some key universe $\mathcal{U}$ to some range $\mathcal{R}$ of hash values is random variable following some distribution over $\mathcal{R}^{\mathcal{U}}$. We say that h is *ε-minwise* or has *bias* ε if for every $S \subseteq \mathcal{U}$ and $x \in \mathcal{U} \setminus S$,

$$\Pr[h(x) \leqslant \min h(S)] \leqslant \frac{1+\varepsilon}{|S|+1} \tag{1}$$

$$\Pr[h(x) < \min h(S)] \geqslant \frac{1-\varepsilon}{|S|+1} \tag{2}$$

From (1) and (2), we easily get for any $A, B \subseteq \mathcal{U}$, that

$$\Pr_h[\min h(A) = \min h(B)] = (1 \pm \varepsilon) \cdot \frac{|A \cap B|}{|A \cup B|} \ .$$

To implement ε-minwise hashing in Wegman and Carter's [23] classic framework of k-independent hash functions $\Theta(\log \frac{1}{\varepsilon})$-independence is both sufficient [11] and necessary [15]. These results are for "worst-case" k-independent hash functions. A much more time-efficient solution is based on simple tabulation hashing of Zobrist [25]. In simple tabulation hashing, the hash value is computed by looking up $c = O(1)$ bitstrings in tables of size $|\mathcal{U}|^{1/c}$ and XORing the results. This is very fast with tables in cache. Pătraşcu and Thorup have shown [16] that simple tabulation hashing, which is not even 4-independent, has bias $\varepsilon = \tilde{O}(1/|S|^{1/c})$. Unfortunately, this bias is useless for small sets S.

In this paper, we consider the twisted tabulation of Pătraşcu and Thorup [17] which was invented to yield Chernoff-style concentration bounds, and high

[2]This is the Iverson bracket notation, where $[P]$ is 1 for a predicate P if P is true and 0 otherwise.

probability amortized performance bounds for linear probing. It is almost as fast as simple tabulation using the same number of lookups but an extra XOR and a shift. We show that with twisted tabulation, the bias is $\varepsilon = \tilde{O}(1/|U|^{1/c})$, which is independent of the set size.

It should be noted, that Thorup [21] recently introduced a double tabulation scheme yielding high independence in $O(1)$ time, hence much faster than using an $\omega(1)$-degree polynomial to get $\omega(1)$-independence and $o(1)$ bias. However, with table size $|U|^{1/c}$, the scheme ends up using at least $7c$ lookups [21, Theorem 1] and 12 times more space, so we expect it to be at least an order of magnitude slower than twisted tabulation[3].

When using minwise for similarity estimation, to reduce variance, we typically want to run q experiments with q independent hash functions $h_1, ..., h_q$, and save the vector of $(\min h_1(A), ..., \min h_q(A))$ as a *sketch* for the set A. We can then estimate the Jaccard similarity as $\sum_{i=1}^{q}[\min h_i(A) = \min h_i(B)]/q$. While q reduces variance, it does not reduce bias, so the bias has to be small for each h_i. This scheme is commonly referred to as $k\times$minwise. Since $\min h_1(A)$ is always compared to $\min h_1(B)$, we say that the samples of the two sketches are *aligned*. A standard alternative[1], called bottom-q, is to just use a single hash function h, and store the q smallest hash values as a set $S(A)$. Estimating the Jaccard-index is then done as $|S(A) \cap S(B) \cap \{q \text{ smallest values of } S(A) \cup S(B)\}|/q$. It turns out that a large q reduces both variance and bias [20]. However, the problem with bottom-q sketches, is that the samples lose their alignment. In applications of large-scale machine learning this alignment is needed in order to efficiently construct a dot-product for use with a linear support vector machine (SVM) [4] such as LIBLINEAR [9] or Pegasos [19]. Using the alignment of $k\times$minwise, it was shown how to construct such a dot-product in [13] based on this scheme. In such applications it is therefore important to have small bias ε. Finally, we note that when $q = 1$, both schemes reduce to basic minwise hashing with the fundamental goal of sampling a single random element from any set with only a small bias, which is exactly the problem addressed in this paper.

2 Preliminaries

Let us briefly review tabulation-based hashing. For both simple and twisted tabulation we are dealing with some universe $\mathcal{U} = \{0, 1, \ldots, u-1\}$ denoted by $[u]$ and wish to hash keys from $[u]$ into some range $\mathcal{R} = [2^r]$. We view a key $x \in [u]$ as a vector of $c > 1$ *characters* from the alphabet $\Sigma = [u^{1/c}]$, i.e. $x = (x_0, \ldots, x_{c-1}) \in \Sigma^c$. We generally assume c to be a small constant (e.g. 4).

[3]The whole area of tabulation hashing is about minimizing the number of lookups, e.g., [12] saves a factor 2 in lookups over [22] for moderate independence.

[4]See `http://en.wikipedia.org/wiki/Support_vector_machine#Linear_SVM`

2.1 Simple Tabulation

In simple tabulation hashing we initialize c tables $h_0, \ldots, h_{c-1} : \Sigma \to \mathcal{R}$ with independent random data. The hash $h(x)$ is then computed as

$$h(x) = \bigoplus_{i \in [c]} h_i[x_i] \ .$$

Here $\oplus$ denotes bit-wise XOR. This is a well-known scheme dating back to [25].

Simple Tabulation is known to be 3-independent, but it was shown in [16] to have much more powerful properties than this would suggest. These properties include fourth moment bounds, Chernoff bounds when distribution balls into many bins and random graph properties necessary in cuckoo hashing. It was also shown that simple tabulation is ε-minwise independent with $\varepsilon = O\left(\frac{\lg^2 n}{n^{1/c}}\right)$.

We will need the following basic lemma regarding simple tabulation ([16, Lemma 2.2]):

Lemma 1. *Suppose we use simple tabulation to hash $n \leqslant m^{1-\varepsilon}$ keys into m bins for some constant $\varepsilon > 0$. For any constant γ, all bins get less than $d = \min\{((1+\gamma)/\varepsilon)^c, 2^{(1+\gamma)/\varepsilon}\}$ keys with probability $\geqslant 1 - m^{-\gamma}$.*

Specifically this implies that if we hash n keys into $m = nu^\varepsilon$ bins, then each bin has $O(1)$ elements with high probability. In this paper "with high probability" (w.h.p.) means with probability $1 - u^{-\gamma}$ for any desired constant $\gamma > 1$.

2.2 Twisted Tabulation

Twisted tabulation hashing is another tabulation-based hash function introduced in [17]. Twisted tabulation can be seen as two independent simple tabulation functions $h^\tau : \Sigma^{c-1} \to \Sigma$ and $h^{\mathcal{S}} : \Sigma^c \to \mathcal{R}$. If we view a key x as the head $head(x) = x_0$ and the tail $tail(x) = (x_1, \ldots, x_{c-1})$, we can define the hash value of twisted tabulation as follows:

$$\begin{aligned} t(x) &= h^\tau(tail(x)) \\ h_{>0}(x) &= \bigoplus_{i=1}^{c-1} h_i^{\mathcal{S}}[x_i] \\ h(x) &= h_{>0}(x) \oplus h_0^{\mathcal{S}}[x_0 \oplus t(x)] \ . \end{aligned}$$

We refer to the value $x_0 \oplus t(x)$ as the *twisted head* of the key x, and define the *twisted group* of a character α to be $G_\alpha = \{x \mid x_0 \oplus t(x) = \alpha\}$. For the keys in G_α, we refer to the XOR with $h_0^{\mathcal{S}}[x_0 \oplus t(x)]$ as the *final (XOR)-shift*, which is common to all keys in G_α. We call $h_{>0}(x)$ the *internal hashing*.

Throughout the proofs we will rely on the independence between h^τ and $h^{\mathcal{S}}$ to fix the hash function in a specific order, i.e. fixing the twisted groups first.

One powerful property of twisted tabulation is that the keys are distributed nicely into the twisted groups. We will use the following lemma from the analysis of twisted tabulation [17, Lemma 2.1]:

Lemma 2. *Consider an arbitrary set S of keys and a constant parameter $\varepsilon > 0$. W.h.p. over the random choice of the twister hash function, h^τ, all twisted groups have size $O(1 + |S|/\Sigma^{1-\varepsilon})$.*

Twisted tabulation hashing also gives good concentration bounds in form of Chernoff-like tail bounds, which is captured by the following lemma, [17, Theorem 1.1].

Lemma 3. *Choose a random twisted tabulation hash function $h : [u] \to [u]$. For each key $x \in [u]$ in the universe, we have an arbitrary* value function *$v_x : [u] \to [0, 1]$ assigning a value $V_x = v_x(h(x)) \in [0, 1]$ to x for each possible hash value. Let $\mu_x = \mathbf{E}_{y\in[u]}[v_x(y)]$ denote the expected value of $v_x(y)$ for uniformly distributed $y \in [u]$. For a fixed set of keys $S \subseteq [u]$, define $V = \sum_{x\in S} V_x$ and $\mu = \sum_{x\in S} \mu_x$. Let γ, c, and ε be constants. Then for any $\mu < \Sigma^{1-\varepsilon}$ and $\delta > 0$ we have:*

$$\Pr[V \geqslant (1+\delta)\mu] \leqslant \left(\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right)^{\Omega(\mu)} + 1/u^\gamma \tag{3}$$

$$\Pr[V \leqslant (1-\delta)\mu] \leqslant \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\Omega(\mu)} + 1/u^\gamma \tag{4}$$

In practice, we can merge h^τ and $h^{\mathcal{S}}$ to a single simple tabulation function $h^\star : \Sigma \to \Sigma \times \mathcal{R}$, but with $h_0^\star : \Sigma \to \mathcal{R}$. This adds $\log \Sigma$ bits to each entry of the tables $h_1^\star, \ldots h_{c-1}^\star$ (in practice we want these to be 32 or 64 bits anyway). See the code in Figure 1 for an implementation of 32-bit keys in C.

```
INT32 TwistedTab32(INT32 x, INT64[4][256] H) {
    INT32 i;
    INT64 h=0;
    INT8 c;
    for (i=0;i<3;i++) {
      c=x;
      h^=H[i][c];
      x = x>> 8;
    }                          // at the end i=3
    c=x^h;                     // extra xor with h
    h^=H[i][c];
    h>>=32;                    // extra shift of h
    return ((INT32) h);
}
```

Figure 1: C-code implementation of twisted tabulation for 32-bit keys assuming a point H to randomly fille storage.

3 Minwise for twisted tabulation

We will now show the following theorem:

Theorem 1. *Twisted tabulation is $O\Big(\frac{\log^2 u}{\Sigma}\Big)$-minwise independent.*

Recall from the definition of ε-minwise, that we are given an input set S of $|S| = n$ keys and a query key $q \in \mathcal{U} \setminus S$. We will denote by Q the twisted group of the query key q. Similarly to the analysis in [16] we assume that the output range is $[0, 1)$. We pick $\ell = \gamma \log u$ and divide the output range into n/ℓ bins. Here γ is chosen such that the number of bins is a power of two and large enough that the following two properties hold.

1. The minimum bin $[0, \ell/n)$ is non-empty with probability $1 - 1/u^2$ by Lemma 3. Here $\mu = O(\log u) < \Sigma^{1-\varepsilon}$.
2. The bins are d-bounded for each twisted group (for some constant d) with probability $1 - 1/u^2$ by Lemma 1. Meaning that for any twisted group G, at most d keys land in each of the n/ℓ bins after the internal hashing is done. This holds because each twisted group has $n/\Sigma^{1-\varepsilon}$ elements w.h.p. by Lemma 2.

Similar to [16], we assume that the hash values are binary fractions of infinite precision so we can ignore collisions. The theorem holds even if we use just $\lg(n\Sigma)$ bits for the representation: Let $\tilde{h}$ be the truncation of h to $\lg(n\Sigma)$ bits. There is only a distinction when $\tilde{h}(q)$ is minimal and there exists some $x \in S$ such that $\tilde{h}(x) = \tilde{h}(q)$. Since the minimum bin is non-empty with probability $1 - 1/u^2$ we can bound the probability of this from above by

$$\Pr\Big[\tilde{h}(q) \leqslant \ell/n \wedge \exists x \in S : \tilde{h}(x) = \tilde{h}(q)\Big] \leqslant \frac{\ell}{n} \cdot \left(n \cdot \frac{1}{n\Sigma}\right) + 1/u^2$$

using 2-independence to conclude that $\{\tilde{h}(q) \leqslant \ell/n\}$ and $\{\tilde{h}(x) = \tilde{h}(q)\}$ are independent.

3.1 Upper bound

To upper bound the probability that $h(q)$ is smaller than $\min h(S)$ it suffices to look at the case when q is in the minimum bin $[0, \ell/n)$, as we have

$$\begin{aligned}\Pr[h(q) < \min h(S)] &\leqslant \Pr[\min h(S) \geqslant \ell/n] + \Pr[h(q) < \min(h(S) \cup \{\ell/n\})] \\ &\leqslant 1/u^2 + \Pr[h(q) < \min(h(S) \cup \{\ell/n\})] \qquad (5)\end{aligned}$$

To bound (5) we will use the same notion of representatives as in [16]: If a non-query twisted group $G_\alpha \neq Q$ has more than one element in some bin, we pick one of these arbitrarily as the representative. Let $R(G_\alpha)$ denote the set of representatives from G_α and let R denote the union of all such sets. We trivially have that $\Pr[h(q) < \min h(S)] \leqslant \Pr[h(q) < \min h(R)]$.

The proof relies on fixing the tables associated with the hash functions h^τ and $h^\mathcal{S}$ in the following order:

1. Grouping into twisted groups is done by fixing h^τ. Each group has $O(1 + n/\Sigma^{1-\varepsilon})$ elements by Lemma 2 w.h.p.

2. The internal hashing of all twisted groups is done by fixing the tables $h_1^{\mathcal{S}}, \ldots, h_{c-1}^{\mathcal{S}}$. This determines the set of representatives R.

3. Having fixed the set R we do the final shifts of the twisted groups G_α by fixing $h_0^{\mathcal{S}}$. We will show that the probability of q having the minimum hash value after these shifts is at most $1/(|R| + 1)$.

   Since $|R|$ is a random variable depending only on the internal hashing and twisted groups, the entire probability is bounded by $\mathbf{E}[1/(|R| + 1)]$.

To see step 3 from above we let $\mathit{Rand}(A)$ be a randomizing function that takes each element in a set A and replaces it with an independent uniformly random number in $[0, 1)$. We will argue that

$$\Pr[h(q) < \min h(R) \cup \{\ell/n\}] \leqslant \Pr[h(q) < \min \mathit{Rand}(R)] = 1/(|R| + 1) \quad (6)$$

To prove (6) fix $h(q) = p < \ell/n$ and consider some twisted group G_α. When doing the final shift of the group we note that each representative $x \in R(G_\alpha)$ is shifted randomly, so $\Pr[h(x) \leqslant p] = p$. However, since the number of bins is a power of two, and each representative in $R(G_\alpha)$ is shifted by the same value, at most one element of $R(G_\alpha)$ can land in the minimum bin. This gives $\Pr[\min h(R(G_\alpha)) \leqslant p] = |R(G_\alpha)|p$. For $\mathit{Rand}(R)$, a union bound gives that $\Pr[\min \mathit{Rand}(R(G_\alpha)) \leqslant p] \leqslant |R(G_\alpha)|p$, implying that

$$\Pr[p < \min(h(R(G_\alpha)) \cup \{\ell/n\})] \leqslant \Pr[p < \min(\mathit{Rand}(R(G_\alpha)) \cup \{\ell/n\})]$$

Because the shifts of different twisted groups are done independently we get

$$\begin{aligned}
\Pr[p < \min(h(R) \cup \{\ell/n\})] &= \prod_{G_\alpha \neq Q} \Pr[p < \min(h(R(G_\alpha)) \cup \{\ell/n\})] \\
&\leqslant \prod_{G_\alpha \neq Q} \Pr[p < \min(\mathit{Rand}(R(G_\alpha)) \cup \{\ell/n\})] \\
&= \Pr[p < \min(\mathit{Rand}(R) \cup \{\ell/n\})] \\
&\leqslant \Pr[p < \min \mathit{Rand}(R)]
\end{aligned}$$

This holds for any value $p < \ell/n$, so it also holds for our random hash value $h(q)$. Therefore

$$\Pr[h(q) < \min(h(R) \cup \{\ell/n\})] \leqslant \Pr[h(q) < \min \mathit{Rand}(R)] \leqslant 1/(|R| + 1)$$

This finishes the proof of (6).

All that remains is to bound the expected value $\mathbf{E}[1/(|R| + 1)]$ and thus the total probability when the internal hashing and twisted groups are random. We will do this using a convexity argument, so we need the following constraints

on the random variable $|R|$: We trivially have $1 \leqslant |R| \leqslant n$. We know that the internal hashing is d-bounded with probability $1 - 1/u^2$, which gives $|R| \geqslant |S \setminus Q|/d \geqslant n/(2d)$. To bound $\mathbf{E}[|R|]$ from below, consider the probability that a key x is *not* a representative. For this to happen x must land in the query group, or another element must land in the same twisted group and bin as x. By 2-independence and a union bound the probability of this event is at most $1/\Sigma + (n-1) \cdot 1/\Sigma \cdot \ell/n = O(\ell/\Sigma)$. The expected number of representatives is therefore

$$\begin{aligned}\mathbf{E}[|R|] &= \sum_{x \in S} \Pr[x \in R] \\ &= \sum_{x \in S} (1 - \Pr[x \notin R]) \\ &\geqslant n \cdot (1 - O(\ell/\Sigma)) \ .\end{aligned}$$

To bound $\mathbf{E}[1/(|R|+1)]$ we introduce a random variable r which maximizes $\mathbf{E}[1/(r+1)]$ while satisfying the constraints of $|R|$ noted above. By convexity of $1/(r+1)$ we get that $\mathbf{E}[1/(r+1)]$ is maximized when r takes the most extreme values. Hence $r = 1$ with probability $1/u^2$, $r = n/(2d)$ with the maximal probability p and $r = n$ with probability $(1 - p - 1/u^2)$. This gives an expected value of

$$\mathbf{E}[r] = 1/u^2 + p \cdot n/(2d) + (1 - p - 1/u^2) \cdot n \ .$$

Thus $p = O(\ell/\Sigma)$ to respect the constraints. To bound $\mathbf{E}[1/(|R|+1)]$ we have

$$\begin{aligned}\mathbf{E}[1/(|R|+1)] &\leqslant \mathbf{E}[1/(r+1)] \\ &\leqslant \frac{1}{2u^2} + \frac{p}{n/(2d)+1} + \frac{1-p-1/u^2}{n+1} \\ &\leqslant \frac{O(p)}{n+1} + \frac{1}{n+1} + O(1/u^2) \\ &= \frac{1}{n+1} \cdot (1 + O(\ell/\Sigma)) \ . \end{aligned} \tag{7}$$

Combining (5), (6) and (7) we get

$$\begin{aligned}\Pr[h(q) < \min h(S)] &\leqslant \Pr[h(q) < \min(h(S) \cup \{\ell/n\})] + O(1/u^2) \\ &\leqslant \Pr[h(q) < \min(h(R) \cup \{\ell/n\})] + O(1/u^2) \\ &\leqslant \mathbf{E}[1/(|R|+1)] + O(1/u^2) \\ &= \frac{1}{n+1} \cdot \left(1 + O\left(\frac{\log u}{\Sigma}\right)\right) \ .\end{aligned}$$

3.2 Lower bound

We have two cases for the lower bound. When $n = O(\log u)$ we observe that the probability of some twisted group having more than one element is bounded from above by $n^2/\Sigma = O(\log^2 u/\Sigma)$ using 2-independence and a union bound.

Since the twisted groups hash independently of each other we have in this case that all elements hash independently. The probability of q getting the smallest hash value is thus at least $1/(n+1) \cdot (1 - O(\log^2 u/\Sigma))$.

When $n = \omega(\log u)$ we again look at the case when q lands in the minimum bin $[0, \ell/n)$. We consider the query group Q separately and thus look at the expression:

$$\begin{aligned}\Pr[h(q) < \min h(S)] &\geqslant \Pr[h(q) < \min(h(S) \cup \{\ell/n\})] \\ &= \Pr[h(q) < \min(h(S \setminus Q) \cup \{\ell/n\})] \\ &\quad - \Pr[\min h(Q) < h(q) < \min(h(S \setminus Q) \cup \{\ell/n\})] \quad . \quad (8)\end{aligned}$$

Furthermore we will assume that all twisted groups have $O(1+n/\Sigma^{1-\varepsilon})$ elements at the cost of a factor $(1 - 1/u^2)$ by Lemma 2. We will subtract this extra term later in (13). Since the twisted groups hash independently we have for a fixed $h(q) = p < \ell/n$ that

$$\Pr[p < \min h(S \setminus Q)] = \prod_{G_\alpha \neq Q} \Pr[p < \min h(G_\alpha)] \quad . \quad (9)$$

We can bound this expression using [16, Lemma 5.1], which states that $1 - pk > (1-p)^{(1+pk)k}$ for $pk \leqslant \sqrt{2} - 1$ and $p \in [0, 1]$. Consider a twisted group G_α and some element $x \in G_\alpha$. We have $\Pr[h(x) < p] = p$ and a union bound gives us that $\Pr[p < \min h(G_\alpha)] \geqslant 1 - p|G_\alpha|$. Since $n = \omega(\log u)$ we have that $p|G_\alpha| \leqslant \ell/n \cdot O(1 + n/\Sigma^{1-\varepsilon}) = o(1)$, so the conditions for the lemma hold. This gives us

$$1 - p|G_\alpha| \geqslant (1-p)^{|G_\alpha|(1+p\cdot|G_\alpha|)} \quad (10)$$

Plugging this into (9) gives

$$\begin{aligned}\Pr[p < \min h(S \setminus Q)] &\geqslant \prod_{G_\alpha \neq Q} (1-p)^{|G_\alpha|(1+p|G_\alpha|)} \\ &\geqslant \prod_{G_\alpha \neq Q} (1-p)^{|G_\alpha|(1+p\cdot 2(|G_\alpha|-1))} \\ &\geqslant (1-p)^m,\end{aligned}$$

with

$$m = n + O(\ell/n) \cdot \sum_{G_\alpha \neq Q} (|G_\alpha| - 1)|G_\alpha| \quad .$$

To bound the entire probability we thus integrate from 0 to ℓ/n:

$$\begin{aligned}\Pr[h(q) \leqslant \min(h(S \setminus Q) \cup \{\ell/n\})] &= \int_0^{\ell/n} \Pr[p < \min h(S \setminus Q)] \, \mathrm{d}p \\ &\geqslant \int_0^{\ell/n} (1-p)^m \mathrm{d}p \\ &\geqslant \frac{1 - (1 - \ell/n)^{m+1}}{m+1} \\ &> 1/(m+1) - 1/(nu) \quad . \quad (11)\end{aligned}$$

Similar to the upper bound m only depends on the twisted groups and their internal hashing, so the entire probability is bounded by $\mathbf{E}[1/(m+1)] - 1/nu \geqslant 1/\mathbf{E}[m+1] - 1/nu$. We note that the sum $\sum_{G_\alpha \neq Q}(|G_\alpha|-1)|G_\alpha|$ counts for each key in a non-query group the number of other elements in its group, so

$$\mathbf{E}\left[\sum_{G_\alpha \neq Q}(|G_\alpha|-1)|G_\alpha|\right] \leqslant n^2/\Sigma \ .$$

The expected value $\mathbf{E}[m+1]$ is therefore bounded by

$$\mathbf{E}[m+1] \leqslant (n+1)\cdot(1+O(\ell/\Sigma)) \ . \tag{12}$$

We can combine this with (11) and get a bound on the first part of (8). We also need to subtract the probability that the keys don't distribute nicely into twisted groups. Doing this we get the following bound:

$$\begin{aligned}\Pr[h(q) \leqslant \min(h(S\setminus Q)\cup\{\ell/n\})] &\geqslant \mathbf{E}[1/(m+1)] - 1/nu - 1/u^2 \\ &\geqslant 1/\mathbf{E}[m+1] - 1/nu - 1/u^2 \\ &\geqslant \frac{1}{(n+1)(1+O(\ell/\Sigma))} - 1/nu - 1/u^2 \\ &\geqslant \frac{1}{n+1}\cdot\left(1 - O\left(\frac{\log u}{\Sigma}\right)\right) \end{aligned} \tag{13}$$

To finish the bound on (8) we need to give an upper bound on

$$\Pr[\min h(Q) < h(q) < \min h(S\setminus Q) \wedge h(q) < \ell/n] \ . \tag{14}$$

To do this we will again consider the set of representatives that we used in the upper bound. We start by fixing the twisted groups. Just like in the upper bound we have w.h.p. that $|R| \geqslant n/(2d)$. We can therefore bound (14) by

$$1/u^2 + \Pr[\min h(Q) < h(q) < \min h(S\setminus Q) \wedge h(q) < \ell/n \wedge |R| \geqslant n/(2d)] \ .$$

We fix $h(q) = p$ for some $p < \ell/n$. Using 2-independence between the fixed query value p and each element of Q we get $\Pr[\min h(Q) < p] \leqslant p|Q|$ and thus

$$\Pr[\min h(Q) < p \wedge |R| \geqslant n/(2d)] \leqslant p|Q| \ . \tag{15}$$

We wish to multiply this by

$$\Pr[p < \min h(S\setminus Q) \mid \min h(Q) < p \wedge |R| \geqslant n/(2d)] \ .$$

For this we use the same approach as for (6). We know that when $p < \ell/n$ we have that $\Pr[p < \min h(R)] \leqslant \Pr[p < \min Rand(R)] = (1-p)^{|R|}$. This holds regardless of the internal hashing so our restriction of $|R| \geqslant n/(2d)$ does not change anything. We now get

$$\Pr[p < \min h(S\setminus Q) \mid \min h(Q) < p \wedge |R| \geqslant n/(2d)] \leqslant (1-p)^{n/(2d)} \ .$$

Multiplying together with (15) we get

$$\begin{aligned}\Pr[\min h(Q) < p < \min h(S \setminus Q) \wedge |R| \geqslant n/(2d)] &\leqslant p|Q|(1-p)^{n/(2d)} \\ &\leqslant p|Q|e^{-pn/(2d)}\end{aligned}$$

for a fixed $p < \ell/n$. To finish the bound we thus integrate from 0 to ℓ/n and get an upper bound on (14):

$$\begin{aligned}&\Pr[\min h(Q) < h(q) < \min h(S \setminus Q) \wedge h(q) < \ell/n] \\ &\quad\leqslant 1/u^2 + \Pr[\min h(Q) < h(q) < \min h(S \setminus Q) \wedge h(q) < \ell/n \wedge |R| \geqslant n/(2d)] \\ &\quad\leqslant 1/u^2 + \int_0^{\ell/n} p|Q|e^{-pn/(2d)}\ dp \\ &\quad= 1/u^2 + O\left(\int_0^{d/n} p|Q|\ dp\right) \\ &\quad= 1/u^2 + O(|Q|/n^2)\ .\end{aligned}$$

We now note that $|Q|$ is a random variable with expected value n/Σ, which gives the final bound on (14) as

$$\begin{aligned}\Pr[\min h(Q) < h(q) < \min h(S \setminus Q) \wedge h(q) < \ell/n] &\leqslant \mathbf{E}\big[O(|Q|/n^2)\big] \\ &= O(1/n\Sigma)\ . \qquad (16)\end{aligned}$$

Combining (8), (13) and (16) gives the desired bound:

$$\begin{aligned}\Pr[h(q) < \min h(S)] &\geqslant \Pr[h(q) < \min h(S \setminus Q) \wedge h(q) < \ell/n] \\ &\quad - \Pr[\min h(Q) < h(q) < \min h(S \setminus Q) \wedge h(q) < \ell/n] \\ &\geqslant \frac{1}{n+1} \cdot \left(1 - O\left(\frac{\log u}{\Sigma}\right)\right) - O\left(\frac{1}{n\Sigma}\right) \\ &= \frac{1}{n+1} \cdot \left(1 - O\left(\frac{\log u}{\Sigma}\right)\right)\end{aligned}$$

References

[1] Yoram Bachrach, Ralf Herbrich, and Ely Porat, *Sketching algorithms for approximating rank correlations in collaborative filtering systems*, Proc. 16th SPIRE, 2009, pp. 344–352.

[2] Yoram Bachrach, Ely Porat, and Jeffrey S. Rosenschein, *Sketching techniques for collaborative filtering*, Proc. 21st IJCAI, 2009, pp. 2016–2021.

[3] Andrei Z. Broder, *On the resemblance and containment of documents*, Proc. Compression and Complexity of Sequences (SEQUENCES), 1997, pp. 21–29.

[4] Andrei Z. Broder, *Identifying and filtering near-duplicate documents*, Proc. 11th CPM, 2000, pp. 1–10.

[5] Andrei Z. Broder, Moses Charikar, Alan M. Frieze, and Michael Mitzenmacher, *Min-wise independent permutations*, Journal of Computer and System Sciences **60** (2000), no. 3, 630–659, See also STOC'98.

[6] Andrei Z. Broder, Steven C. Glassman, Mark S. Manasse, and Geoffrey Zweig, *Syntactic clustering of the web*, Computer Networks **29** (1997), 1157–1166.

[7] Edith Cohen, Mayur Datar, Shinji Fujiwara, Aristides Gionis, Piotr Indyk, Rajeev Motwani, Jeffrey D. Ullman, and Cheng Yang, *Finding interesting associations without support pruning*, IEEE Trans. Knowl. Data Eng. **13** (2001), no. 1, 64–78.

[8] Mayur Datar and S. Muthukrishnan, *Estimating rarity and similarity over data stream windows*, Proc. 10th ESA, 2002, pp. 323–334.

[9] Rong-En Fan, Kai-Wei Chang, Cho-Jui Hsieh, Xiang-Rui Wang, and Chih-Jen Lin, *LIBLINEAR: A library for large linear classification*, Journal of Machine Learning Research **9** (2008), 1871–1874.

[10] Monika Rauch Henzinger, *Finding near-duplicate web pages: a large-scale evaluation of algorithms*, Proc. ACM SIGIR, 2006, pp. 284–291.

[11] Piotr Indyk, *A small approximately min-wise independent family of hash functions*, Journal of Algorithms **38** (2001), no. 1, 84–90, See also SODA'99.

[12] Toryn Qwyllyn Klassen and Philipp Woelfel, *Independence of tabulation-based hash classes*, Proc. 10th Latin American Theoretical Informatics (LATIN), 2012, pp. 506–517.

[13] Ping Li, Anshumali Shrivastava, Joshua L. Moore, and Arnd Christian König, *Hashing algorithms for large-scale learning*, Advances in Neural Information Processing Systems, 2011, pp. 2672–2680.

[14] Gurmeet Singh Manku, Arvind Jain, and Anish Das Sarma, *Detecting near-duplicates for web crawling*, Proc. 10th WWW, 2007, pp. 141–150.

[15] Mihai Pătraşcu and Mikkel Thorup, *On the k-independence required by linear probing and minwise independence*, Proc. 37th International Colloquium on Automata, Languages and Programming (ICALP), 2010, pp. 715–726.

[16] ———, *The power of simple tabulation-based hashing*, Journal of the ACM **59** (2012), no. 3, Article 14, Announced at STOC'11.

[17] Mihai Pătraşcu and Mikkel Thorup, *Twisted tabulation hashing*, Proc. 24th ACM/SIAM Symposium on Discrete Algorithms (SODA), 2013, pp. 209–228.

[18] Saul Schleimer, Daniel Shawcross Wilkerson, and Alexander Aiken, *Winnowing: Local algorithms for document fingerprinting*, Proc. SIGMOD, 2003, pp. 76–85.

[19] Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro, *Pegasos: Primal estimated sub-gradient solver for svm*, Proceedings of the 24th International Conference on Machine Learning, ICML '07, 2007, pp. 807–814.

[20] Mikkel Thorup, *Bottom-k and priority sampling, set similarity and subset sums with minimal independence*, Proc. 45th ACM Symposium on Theory of Computing (STOC), 2013.

[21] ———, *Simple tabulation, fast expanders, double tabulation, and high independence*, FOCS, 2013, pp. 90–99.

[22] Mikkel Thorup and Yin Zhang, *Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation*, SIAM Journal on Computing **41** (2012), no. 2, 293–331, Announced at SODA'04 and ALENEX'10.

[23] Mark N. Wegman and Larry Carter, *New classes and applications of hash functions*, Journal of Computer and System Sciences **22** (1981), no. 3, 265–279, See also FOCS'79.

[24] Hui Yang and James P. Callan, *Near-duplicate detection by instance-level constrained clustering*, Proc. 29th SIGIR, 2006, pp. 421–428.

[25] Albert Lindsey Zobrist, *A new hashing method with application for game playing*, Tech. Report 88, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 1970.